A Case-Based Approach to Open-Ended Collective Agreement with Rational Ignorance

Sergio Manzano[1,2], Santiago Ontañón[1], and Enric Plaza[1]

[1] IIIA-CSIC, Artificial Intelligence Research Institute (Spanish Scientific Research Council), 08193 Bellaterra, Catalonia, Spain
[2] Universitat Autonòma Barcelona, Bellaterra, 08193 Catalonia, Spain
{sergio,santi,enric}@iiia.csic.es

Abstract. In this paper we focus on how to use CBR for making collective decisions in groups of agents. Moreover, we show that using CBR allows us to dispense with standard but unrealistic assumptions taken in these kind of tasks. Typically, social choice studies voting methods but assumes complete knowledge over all possible alternatives. We present a more general scenario called *open-ended deliberative agreement with rational ignorance (ODARI)*, and show how can CBR be used to deal with rational ignorance. We will apply this approach to the *Banquet Agreement* scenario, where two agents deliberate and jointly agree on a two course meal. Rational ignorance makes sense in this scenario, since it would be unreasonable for the agents to know all the alternatives. Unknown alternatives, as well as a strategy to increase chances of reaching an agreement, are problems addressed using case-based methods.

1 Introduction

Case-based reasoning (CBR) has been applied to a wide number of of real life tasks, and one feature that stands out in comparison to other AI techniques is its resilience and robustness in the presence of incomplete knowledge. This capability of dealing with incomplete knowledge, by analyzing and exploiting the implicit knowledge in a case base, is a core idea in CBR, and a main factor in being able to perform well in real-life applications in which standard oversimplifying assumptions (like having complete knowledge) can not be held.

In this paper we focus on how to use CBR for making collective decisions in groups of agents. Moreover, we show that using CBR allows us to dispense with standard but unrealistic assumptions taken in these kind of tasks. Social choice is the theoretical study of methods for aggregating individual preferences (or utilities) into collective decisions —and how to evaluate adequacy of these collective outcomes (welfare). Typically, social choice studies voting methods and their properties as a means to aggregate individual preferences into a collective outcome. However, current social choice approaches make strong assumptions that are tantamount to require complete knowledge to the individual agents participating in a collective decision (see Section 2).

A. Ram and N. Wiratunga (Eds.): ICCBR 2011, LNAI 6880, pp. 107–121, 2011.

We claim that these assumptions are too strong for Artificial Intelligence applications, specially in multiagent systems approaches, and that we should be focusing on group decisions where individuals have partial knowledge of the decision domain. Specifically, we propose that CBR offers a practical and natural way to allow individual agents with partial knowledge to achieve collective decisions that are reasonable and satisficing[1].

Let us consider a group decision scenario in a context related to the ICCBR Computer Cooking Competition. While the main challenges of the Computer Cooking Competition are retrieving and adapting a recipe given some query, we propose to focus on the task we call *Banquet Agreement scenario*[2]. Let us consider a group of 2 or more individuals that have to decide on the specific dishes to be served in a banquet: these individuals may be the Chairs of a conference deciding on the Conference Banquet, or the members of two families for a marriage banquet, or a couple inviting a large group of friends to dinner.

This scenario allows us to consider the implications of assuming complete or incomplete knowledge. Classical social choice assumes complete knowledge: every individual *knows* the utility value of every alternative —usually this is phrased as every individual having a utility function over the set of alternatives. However, in the cooking domain, the set of alternatives are the different dishes or recipes, for which there are thousands. Clearly, assuming that an individual has a utility value for *every and all recipes* (complete knowledge) is unfeasible. Moreover, in a group decision, the set of known alternatives (i.e. those for which the utility value is known) to each individual may differ, and the set of alternatives known to *all individuals* in the group may be small or empty.

Thus, in any realistic scenario, specially in group decision, assuming complete knowledge is too restrictive and we have to deal with incomplete knowledge. This means, in the Banquet Agreement scenario, that each individual may know a specific subset of recipes (and thus have utility value for them), but not the rest. In fact, some models of economics recognize this possibility, which is called *rational ignorance*. The notion of rational ignorance means that, in those situations where the cost of acquiring information is greater than the benefits to be derived from the information, it is rational to be ignorant. Thus, in the cooking domain, it is rational to be ignorant — otherwise the time and cost involved in finding all possible recipes, tasting them, and acquiring an individual utility value would be a (maybe pleasurable but) daunting task.

Our approach is based on considering the set of alternatives under discussion to be open ended —instead of being fixed beforehand by an external entity, the alternatives are introduced in the discussion by the individual agents. This

[1] The word *satisfice* was proposed by Herbert Simon in 1956 as an alternative desideratum for AI tasks, in contrast with the more classical word *maximize*; Simon championed the notion of *bounded rationality* as a more realistic approach to rationality that takes into account different kinds of cognitive limitations.

[2] We will consider group decisions as a form of group agreement. In our approach, negotiation, mediation, social choice, etc. are different kinds of processes whose aim is to reach agreements between 2 or more agents. In what follows, we will speak of group decision or agreement interchangeably.

process by which agents try to reach an agreement (and in which the set of alternatives is expanded) will be called *deliberative agreement*[3]. For instance, in the Banquet Agreement scenario one agent may know about one risotto recipe (e.g. mushroom risotto), but once another agent proposes as a candidate agreement another risotto recipe (e.g. fresh asparagus risotto), the set of what we call *public alternatives* is increased. Then, the first agent can deal with one more alternative, and although rational ignorance implies that the agent has no utility value for an unknown alternative, case-based reasoning can be used to estimate utility values of unknown public alternatives by such an agent. Therefore, CBR can be used to deal with rational ignorance (incomplete information) based on what is known by each agent. This capability allows us to deal with more realistic scenarios with an open-ended set of alternatives.

Moreover, we will show that CBR can also be used during deliberative agreement process to reach faster a group agreement. A CBR agent can make a case-based model of another agent using the proposals offered by that agent. In the Banquet Agreement scenario, for instance, if one agent proposes the alternative fresh asparagus risotto, another agent may safely infer that the agent would also like similar recipes, and use CBR to try to propose a preferred recipe (e.g. risotto with "fava" beans) that is also similar to recipes proposed by another agent. Although our approach is valid for n agents, in most of the paper we will focus on the 2 agent scenario for sake of expository clarity

In this paper, we will present an approach to group agreement in multiagent systems in which CBR is used in situations where rational ignorance applies. Case-based reasoning will be shown to deal with incomplete knowledge of individuals and, moreover, support satisficing behavior in combinatorial domains in multi-issue group agreement. The next section introduces the classical notions of single-issue and multi-issue collective decisions, while the rest of the paper presents our approach to open-ended group agreement in Section 3, and the CBR strategy to deal with group agreement and rational ignorance in Section 5. An experimental evaluation in the cooking domain is shown in Section 7, and the paper comes to an end with sections on related work and conclusions.

2 Background

Social choice theory [3,6,1] is a mathematical theory of collective decision making, which is concerned about how groups actually do make decisions, focusing on methods to aggregate individual preferences into a collective decision or choice. A social choice problem consists of a number of individuals that have preferences over a set of alternatives $D^i = x^1, \ldots, x^k$ on an issue X^i. A preference aggregation method, typically a voting method, aggregates the preferences over alternatives, ranks those alternatives in a global ordering, and determines the

[3] While classical social choice focuses on voting methods, current research in deliberative democracy (and deliberation in multimember bodies in general) show that in real life situations it is useful having a first stage of deliberation before the stage of voting on a collective outcome [4].

winning alternative for the issue (the collective decision). Voting methods include the majority rule, approval voting, Borda count, the Condorcet method, etc. Social choice theory studies the properties satisfied or not by different aggregation methods. However, they share some basic assumptions: agents are supposed to have perfect knowledge: they know all possible alternatives $D^i = x^1 \ldots x^k$ on the issue at hand X^i and know which ones are preferred over the others.

When the decision is more complex, such decision is modeled as a set of issues, each one with a number of alternatives. Thus, multi-issue social choice consists of a set of issues $X = \{X^1, \ldots, X^m\}$ and for each issue X^k there might be a set of alternatives x to choose from, in a domain D^k. The number of possible decisions is now much greater: all possible combinations in $\Omega = D^1 \times \ldots \times D^m$. Assuming the issues are independent (i.e. preferences are separable) is a common simplification that allows issue-by-issue voting to achieve a group multi-issue decision. However, the individuals' preferences are not necessarily separable, since the issues may be interdependent, i.e. that an individual's preference for one issue, may depend on the alternative taken for another issue. For instance, in the Banquet Agreement scenario, a group has to decide on a two-course meal. Clearly, an individual's preference on the main course may depend on the alternative taken as a starter, e.g. having "arròs rossejat" as main course rules out having a rice salad as a starter, since both alternatives' main ingredient is rice.

Thus, in multi-issue social choice with dependences between issues, the issue-by-issue voting method may lead to suboptimal results. For instance, in the Banquet Agreement scenario, a subgroup constituting a majority of individuals may select "rice salad" as a starter while another sub group constitute a majority selecting "arròs rossejat" as main course even when no individual votes for having both rice salad and "arròs rossejat" together. The approach we will take is that multi-issue social choice problem is composed of the possible combinations in $D^1 \times \ldots \times D^m$ and a set of constraints C over them in such a way that we define a set of valid combinations $\Omega \subset D^1 \times \ldots \times D^m$.

A different but related approach to collective decision making is of that of *deliberation* in the study of deliberative democracy [2]. Deliberative democracy contends that collective decision processes should not just aggregate individual preferences but help shaping those individual preferences. Therefore, deliberative democracy encourages the individuals to deliberate about which alternative is to be preferred for an issue, in an open dialogue with one another, before voting. During the public discussion, the individuals may change their preferences since they can acquire new alternatives for the issues at hand, and other relevant information, such as the concerns of the other individuals. Thus, votes and preferences should emerge from processes of deliberation, since then individuals are able to make a more informed decision.

Given that group decisions are not necessarily limited to a process of choosing among given alternatives, but also a process of generating new alternatives (brainstorming), an argument in favor of discussing publicly before voting, is that the limitations due to bounded rationality might be alleviated, since deliberation can provide more creative outcomes [5]. In our approach, a process of deliberation is required for allowing the individual agents to introduce new alternatives to an issue —deliberation makes possible to take collective decisions within a context of rational ignorance (incomplete knowledge).

From the point of view of CBR, this approach continues the work on multiagent case-based reasoning that previously focused on classification tasks [9,7,8]. In this paper, multiagent CBR is used not to solve a problem (find a correct outcome) but reach an agreement (find a decision with high group welfare). Moreover, the CBR approach (together with deliberation democracy approach) are used to develop a more realistic framework for group choice that embrace openness and incompleteness of knowledge.

3 Open-Ended Deliberative Agreement

This section introduces the open deliberative agreement with rational ignorance (ODARI) framework, in which a group of agents deliberates on a set of issues and their alternatives in order to reach an agreement about the alternatives (rational ignorance). Although ODARI is defined for a group on n agents, it is easier to explain the case where there are 2 agents deliberating on an agreement, and we will use a set of 2 agents $\mathcal{A} = \{A_1, A_2\}$ in our exposition here. How CBR is used to deal in the ODARI framework is explained later in Section 5.

The ODARI framework is a model of group decisions in multiagent systems with the following properties:

Open-Endedness. The first feature in ODARI is that it is a multi-issue group decision problem $X = \{X^1, \ldots, X^m\}$ where each issue X^k is open-ended, i.e. the set of alternatives D^k is not fixed and known by all individual agents (rational ignorance).

Deliberation. A second feature is that new alternatives to issues can be introduced during a process of deliberation to reach an agreement.

Interrelated Issues. A third feature is that there is a set of constraints C that determine a subset of valid combinations $\Omega \subset D^1 \times \ldots \times D^m$

Time-sensitiveness. Finally, since the space of valid combinations Ω can be very large, we assume there is a finite time limit that precludes the exploration of all valid combinations in Ω during deliberative agreement process.

We contend that these four features makes the ODARI framework closer to realistic scenarios of group decisions in multiagent systems.

An issue is open-ended when the set of alternatives it may take increases monotonically over time. An agent A_i has a limited experience in each issue X^k, and knows a subset of all alternatives that may exist in the world, which we will denote as D_i^k. The *agreement space* of an agent A_i based on its initial knowledge of the world is $\Omega_i \subset D_i^1 \times \ldots \times D_i^m$, the set of combinations of the individually known alternatives that satisfy the constraints in C.

Thus, ODARI assumes that an alternative $x \in D_i^k$ may be a known alternative to an agent, but unknown to another, and every agent has a utility value (or preference) for each known alternative to him. In our CBR approach, discussed later in Section 5, this experience-based knowledge will be determined by the individual case base of each agent. Moreover, on the Banquet Agreement scenario, the case base is composed of the cooking recipes known to an agent.

The preferences of an agent A_i for each issue X^k will be expressed by a utility function: $U_i^k : D_i^k \rightarrow [0,1]$. However, A_i has no utility value for unknown alternatives (those not in D_i^k). The domain of alternatives for an issue X^k given a group of agents $A_1, \ldots, A_n$ is $D^k = \bigcup_{i=1\ldots n} D_i^k$, where every agent knows a subset D_i^k. The space of possible combinations is $D^1 \times \ldots \times D^m$, but given a set of constraints C, the space of possible agreements is the subset $\Omega \subset D^1 \times \ldots \times D^m$ of combinations satisfying C. Individual agents have no immediate access to this larger space, since they are working only in a subspace $\Omega_i \subset \Omega$.

During the deliberation process the agents may come to know new alternatives for the issues at hand as they are included in proposals made by other agents. If the agents use this new alternatives the space of possible agreements that may be proposed increases accordingly. Moreover, since this alternatives are presented in a public space (all participating agents have access to all information flow during deliberation), all agents become aware of the new alternatives.

Let l Π^k be the alternatives for the issue X^k made public at some moment in time during deliberation; the agreement space of public alternatives is then $\Omega^P = \Pi^1 \times \ldots \times \Pi^m$. In this situation, Ω^P is the "common knowledge" of the agent group, but the agents still need some way to integrate the unknown alternatives in their preference structure, i.e. the set of "new" alternatives $N_i^k = \Pi^k - D_i^k$ for every issue X^k. Section 5 explains how CBR is used to deal with these new alternatives. Notice, however, that an agent using public alternatives can now generate a larger set of proposals: for each issue X^k the agent can choose from a larger set of alternatives, namely $D_i^k \cup N_i^k$. Therefore, the set of agreements that can be proposed by an agent is now $\Omega_i^P \subset D_i^1 \cup N_i^1 \times \ldots \times D_i^m \cup N_i^m$ (the set of combinations satisfying the set of constraints C).

4 Deliberation Process

The deliberation process among agents is an interaction protocol on which agents propose and accept (or reject) possible agreements of the form $(x^1, \ldots, x^m)$, i.e. assigning one alternative to each of the m issues involved in the deliberative agreement task. For instance, in the Banquet Agreement scenario, with two main courses, an example proposal may be (Xató, arròs-rossejat), where the dishes specify their ingredients: Xató is a endive salad with cot, olives, etc. and a hot sauce, and arròs-rossejat is a fishermen's fried noodles dish with aioli. A *valid* proposal is a combination of alternatives that satisfies the set of constraints C.

In this section we will present the interaction protocol DAP2 for 2 agents participating in the deliberation; extending DAP2 to n agents is possible but the 2 agent scenario is easier to understand. DAP2 allows a group of two agents $\mathcal{A} = \{A_1, A_2\}$ making proposals until one agent accepts a proposal made by the other agent (which becomes the agreement); each proposal is made in a new round of the protocol. A maximum number of rounds M establishes a deadline for reaching an agreement.

During deliberation the agents interchange the following types of messages:

- *propose*(A_i, ω, t): where an agent A_i proposes a valid combination of alternatives ω at the round t; moreover ω has to be a new combination (i.e. ω has never been proposed before).
- *accept*(A_i, ω, t): where an agent A_i accepts the valid combination of alternative ω at the round t; the proposal ω has been proposed previously by the other agent A_j but need not be the last proposal of A_j.

The DAP2 protocol starts at the round $t = 0$ with the token randomly assigned to an agent:

1. The agent A_i who has the token can act in different ways:
   - A_i accepts a previous proposal ω of the other agent A_j sending it message *accept*(A_i, ω, t) and then the protocol terminates with agreement ω.
   - A_i makes a new proposal ω sending *propose*(A_i, ω, t) to the agent A_j; if A_i is unable to find a new proposal an *abstain*(A_i, t) is sent. The token passes to the other agent A_j, and the protocol moves to the step 2.
2. If the deadline M is reached or none of the agents made a proposal in the previous two rounds, the protocol terminates without an agreement. Otherwise a a new round $t + 1$ starts and the protocol moves to the step 1.

For using the DAP2 protocol, agents just need a decision policy that allows them to decide how to act in the protocol (when to accept, and when to make new proposals). We present such policies in the next section.

5 CBR Agents

This section presents how case-based reasoning can address, in a natural way, the challenges associated with a more realistic scenario for collective decision (essentially knowledge incompleteness) in the context of deliberative agreement, where the deliberation process allows agents to acquire new and unknown alternatives. Let use define some auxiliary notions before presenting CBR-based decision policies to be used with the DAP2 protocol.

Open Minded Strategy. We will define two different strategies, open-minded and narrow-minded reasoning strategies. An *open-minded* agent will consider acceptable agreements containing alternatives that are new and unknown for that agent. The *narrow-minded* reasoning strategy, on the other hand, will not consider acceptable any agreement containing unknown alternatives. In the Banquet Agreement scenario, for instance, an agent may like paella but does not know arròs rossejat: the narrow-minded would not accept any agreement with arròs rossejat (even when it is very similar to paella), while the open-minded agent will consider the similarity and may accept agreements with arròs rossejat. Clearly, the open-minded strategy allows a larger space of possible agreements, while the narrow-minded strategy constrains the space of possible agreements to those containing alternatives known to all agents. For instance, between two agents, a narrow-minded strategy has an agreement space $\Omega_1 \cap \Omega_2$, while an open-minded strategy has an agreement space $\Omega_1^P \cup \Omega_2^P$.

Furthermore, an open-minded agent can also use new unknown alternatives when proposing an agreement. For instance, an agent that proposed (green-salad, paella) has later received a proposal containing arròs rossejat; since this means that the other agent likes arròs rossejat and it is similar to paella, the agent can now make a new proposal with more chances to succeed: (green-salad, arròs rossejat). Thus, as we will see in the experimental evaluation section, open-mindedness helps in reaching an agreement faster (while maintaining high levels of satisfaction) by allowing to propose agreements closer to the preferences of the other agent. Narrow-mindedness, on the other hand, risks running out of time without finding a common agreement.

Multi-issue Deliberation. Generating and evaluating proposals of agreement is much more complex in multi-issue group agreement than in single-issue group agreement. Moreover, constraints over the combination of alternatives make infeasible estimating the utility of alternatives in isolation; thus, utility will be measured by a function over possible agreements $U_i : \Omega \rightarrow [0, 1]$. Therefore the CBR agents have to reason about valid combinations of alternatives, but the combinatorial nature of this process together with time-sensitivity makes impossible an approach based on maximization: an approach based on satisficing is needed, where an agent A_i accepts an agreement ω if it is satisfactory to A_i. Later in Section 6 we will formalize this idea with the notion of the aspiration level of an agent, such that when a proposed agreement's surpasses the aspiration level the agent accepts that agreement.

In general, the agents should be able to evaluate any agreement in Ω; thus, for each issue X^k, an agent requires a way of evaluating the utility degree of every alternative in D^k. In the classical approach studied in social choice, since the alternatives are just a set of identifiers, without an intrinsic structure, all knowledge resides in the utility function. However, in ODARI, there are unknown alternatives for any agent A_i —i.e. A_i does not have utility degree for some of alternatives that an issue may take. For this reason, we assume that each CBR agent A_i has a similarity measure over the alternatives of an issue $s_i^k : D^k \times D^k \rightarrow [0, 1]$. This new assumption involves access to some characterization of the alternatives, and the similarity measure works upon that characterization. Consequently, we assume (1) that the alternatives have some characterization or description in some language, and (2) that a similarity measure can be defined on the space of descriptions of alternatives. Clearly, the similarity among alternatives is domain specific; Section 7 presents the similarity we have used in the Banquet Agreement scenario.

Issue Case Base. In our approach, a CBR agent A_i has a case base C_i^k for each issue X^k, where a case $c \in C_i^k$ is a pair $c = \langle x, u \rangle$ such that x is a known alternative to A_i for an issue $X^k$and u is the utility degree of x to A_i, i.e. $c.u = U_i^k(x)$.

Next, using these issue case bases, we will be able to define a function $\overline{U}_i^k$ that estimates the utility of unknown alternatives for each issue. Using $\overline{U}_i^k$ an agent A_i is able to evaluate each possible agreement in Ω, and thus, is able both to accept a proposal, even if it has unknown alternatives, and to use unknown alternatives in generating proposals.

Utility of an Unknown Alternative. The similarity measure allows us to estimate the utility of an unknown alternative x^r of an issue X^r, by using a k-nearest neighbor method which is calculated as the weighted sum of utilities of the k most similar cases in C_i^r:

$$\overline{U}_i^r(x^r) = \frac{\sum_{c \in \mathbb{K}} s_i^r(c.x, x^r) \times c.u}{\sum_{c \in \mathbb{K}} s_i^r(c.x, x^r)}$$

where $\mathbb{K}$ is the set of k most similar cases.

Utility of an Alternative. An agent A_i either knows the utility of an alternative or can estimate it for unknown alternatives.

$$\mathbf{U}_i^k(x) = \begin{cases} U_i^k(x) \text{ if } x \in D_i^k \\ \overline{U}_i^k(x) \text{ otherwise} \end{cases}$$

Utility of an agreement. The utility of an agreement ω (a valid combination of alternatives for m issues) is $U_i(\omega) = \frac{1}{m} \sum_{1 \le k \le m} \mathbf{U}_i^k(x^k)$.

Modeling Agent Preferences. An agent does not have any a priori information about the other agent's preferences. During the deliberation process, an agent does not know, for any ω in the agreement space that has never been proposed, whether ω could be satisfactory to the other agent. However, an agent A_i may exploit the information that emerges during the deliberation, in order to acquire clues about the other agents' preferences. Specifically, when an agent A_i has the token, A_i is aware that the proposals made by A_j up to that point are satisfactory to A_j, because as soon as A_j has proposed an agreement, he commits to accept that agreement. Additionally, A_i is aware that no agreement it has proposed up to that point has been satisfactory to A_j, otherwise A_j would have accepted one of them.

For this reason, when an agent A_i makes a proposal that is similar to proposals made by the agent A_j in previous rounds, A_i increases the likelihood of it being accepted by A_j, since the proposals made by A_j are satisfactory to A_j.

Moreover, the agent A_i may exploit the information about the proposals it made that were rejected by A_j. Here, the intuition is that those proposals give some information to A_i of what kind of proposal are *not* satisfactory to A_j. This way, the more similar a possible agreement to the proposals made by A_i, the more likelihood it will not be satisfactory to A_j.

Thus, an agent A_i may use the similarity between proposals to estimate the proposals' likelihood of being satisfactory (or not) to another agent A_j, based in the proposals previously made by both agents. In this sense, the set of proposals have been made during a deliberation process are treated as a "case base" that models A_j's preferences.

Proposal Case Base. Every agent A_i in ODARI has a proposal case base C_i^P such that, for each proposal ω made by any agent a in the deliberation (including itself), there is a case $\langle \omega, a \rangle \in C_i^P$.

Proposal Similarity. An agent A_i has a similarity function $S_i : \Omega \times \Omega \rightarrow [0, 1]$, which expresses the similarity between two proposals (two valid combination of alternatives):

$$S_i(\{x_1^1, \ldots, x_1^m\}, \{x_2^1, \ldots, x_2^m\}) = \frac{1}{m} \sum_{1 \leq k \leq m} s_i^k(x_1^k, x_2^k)$$

Let $M_i = \{\omega \in C_i^P | c.a = A_i\}$ be the set of proposals made by A_i to the agent A_j. Notice that proposals in M_i have not been accepted by A_j, since those proposals are not satisfactory to A_j; thus, other proposals similar to M_i are likely to be unsatisfactory to A_j. Let $R_i^j = \{\omega \in C_i^P | c.a = A_j\}$ be the proposals received by A_i from A_j, i.e. the set of proposals that are known to be satisfactory to A_j. Therefore, proposals similar to those in R_i^j are more likely to be satisfactory A_j. Thus, the likelihood of a proposal to be accepted by A_j increases if it is similar to proposals in R_i^j and decreases if it is similar to proposal in M_i.

Proposal Acceptance Likelihood. Following these two heuristic criteria, based in the proposals made by the agents, we will define a function E_i allowing A_i to estimate the likelihood of a new agreement proposal ω' being accepted by another agent as follows:

$$E_i(A_j, \omega') = \frac{1}{2}\Big(\max_{\omega \in R_i^j} S_i(\omega', \omega)\Big) + \frac{1}{2}\Big(\min_{\omega \in M_i} (1 - S_i(\omega', \omega))\Big)$$

6 Proposal Generation and Aceptance

A CBR agent A_i, in the ODARI framework, will propose agreements taking into account (1) the proposals' utility degree to A_i (possibly estimated with a similarity) and (2) the likelihood of the proposals being accepted by the other agent A_j. If A_i follows the open-minded strategy, the agreement proposal has to be selected from the space of possible agreements Ω_i^P. Since a proposal cannot be repeated, however, the space of possible new proposals is $\Omega_i' = \Omega_i^P - P$, where P is the set of agreements already proposed. We define first how the agreement to be proposed is selected by an agent.

Proposal Selection Heuristic. The candidate agreements Ω_i' will be evaluated by a heuristic $H_i : \Omega_i' \rightarrow [0, 1]$ that combines the utility for the proposing agent and the likelihood to be accepted by the other agent, as follows:

$H_i(\omega) = (1 - \alpha) \times U_i(\omega) + \alpha \times E_i(A_j, \omega)$.

where $\alpha \in [0, 1]$ is the weight given to the acceptance likelihood estimated for the other agent. The selection of the agreement to be proposed is simply $\mathbf{H}_i(\Omega_i') = argmax_{\omega \in \Omega_i'} H_i(\omega)$ (i.e. the agreement with highest H_i value).

Now we turn to the issue of an agent deciding to accept or not a proposed agreement, based on the notion of aspiration level. The main idea is twofold: (1)

an agent A_i accepts an agreement ω when its utility $U_i(\omega)$ is above its aspiration level ∂_i, and (2) the aspiration level ∂_i decreases during the deliberation process.

Aspiration Level. At any moment in time, an agent A_i has proposed M_i agreements. Since all agreements in M_i are satisfactory to A_i, let us take the one with minimum utility $\omega_i^* = argmin_{\omega \in M_i} U_i(\omega)$; therefore $\partial_i = U_i(\omega_i^*)$ — i.e. that the aspiration level is $U_i(\omega_i^*)$, since agent A_i is already proposing an agreement with that utility degree. Any agreement proposal ω that agent A_i receives whose utility for A_i is equal or greater than the aspiration level should be accepted by A_i (since A_i is already proposing agreement with that degree of utility). Notice that ∂_i will decrease monotonically with time, since the set M_i increases with new proposed agreements and no proposed agreement can be withdrawn according to the protocol DAP2.

Decision Policy. Whenever an agent A_i owns the token in protocol DAP2, the decision policy of an agent A_i decides either to accept an agreement proposed by A_j or to propose a new agreement. Now, agent A_i has an aspiration level ∂_i and a new agreement to propose, namely $\mathbf{H}_i(\Omega'_i)$. Let $\omega^k = argmax_{\omega \in R_i(j)} U_i(\omega)$ be the best proposal received by A_i from A_j, then the decision policy is:

1. if $U_i(\omega^k) \geq min(\partial_i, U(\mathbf{H}_i(\Omega'_i)))$ then Accept ω_k
2. otherwise Propose $\mathbf{H}_i(\Omega'_i)$.

i.e. an agent A_i will accept the best proposal received if it has an utility better or equal than the aspiration level or the utility of the next agreement that A_i intends to propose next; otherwise A_i proposes that agreement.

7 Banquet Agreement Scenario

In this section, we experimentally evaluate our approach on the two-issues Banquet Agreement scenario, based on the recipes database of the 2010 Computer Cooking Contest (CCC). Specifically, these experiments involve two agents A_1 and A_2, that will engage in the deliberative process to reach an agreement on a two-course meal, i.e. agreeing on a specific recipe for the starter (issue one) and for the main course (issue two). The agreement must satisfy the following constraint: no main ingredient may be used in both recipes[4]. The experimental database R consists of 600 recipes (with 380 different ingredients) from the CCC, randomly split into two sets: R^1 and R^2 of starters and main courses.

The domain-specific similarity functions among alternatives of an issue is here a similarity s_1 over recipes in R^1 and s_2 over recipes in R^2. Both s_1 and s_2 are based on the Jaccard similarity

$$s(x_1, x_2) = \frac{|Ing(x_1) \cap Ing(x_2)\ |}{|Ing(x_1) \cup Ing(x_2)\ |}$$

where $Ing(x)$ is the set of main ingredients in recipe x.

[4] Main ingredients like rice or potato cannot be repeated, but secondary ingredients like oil or salt can be used in both course's recipes.

The agents' preferences are built by the experiment designer in the following way. Each agent has a profiler-creation function that randomly assigns a utility value in $[0, 1]$ to each ingredient at each run of the experiment. From the utility of ingredients the utility of a recipe x is computed as the normalized sum of the utility of the ingredients in x. The recipes known to a agent constitute a case base where there is a case $\langle x, u\rangle$ for each known recipe x and its utility u. Notice that (1) the agent does not know the utility of ingredients, only the global utility of recipes, so it is unable to ascertain the utility of unknown recipes from their ingredients, and (2) people usually have clear utilities for courses rather than ingredients (although some ingredients may be a no-no), using the hidden profiler-creation function is just a convenient way to generate a large number of profiles for experimentation.

Moreover, random profiles constitute a rather worst case scenario, where any commonality of tastes among two artificial agents might be much lower than other scenarios where participants may share some tastes. Finally, notice that an agent only knows the utility for the recipes in its case base and is ignorant of other recipes contained in other case bases (except for the recipes both agents know, which will be a smaller subset in the experiments).

7.1 Experiments

In our experiments, we have set the maximum number of rounds for DAP2 to $M = 150$, unless specified otherwise, and we have used $k = 5$ for the k-nearest neighbor method $\overline{U}_i(x)$ to estimate the utility degree of unknown alternatives. Given the maximum number of M rounds, if this maximum is reached the protocol terminates without agreement. The experiments, unless specified otherwise, assign 400 recipes to each agent, 200 of which are shared by both agents.

In order to evaluate the quality of the agreements reached by the agents, we need to define a function assessing the degree of "goodness" of these agreements for the group. These functions are called *social welfare functions*, and there are several that are defined in the literature, depending on the criteria of what does it mean for a agreement to be good for the group. The utilitarian welfare $W_U(\omega) = \frac{1}{2}(U_1(\omega)+U_2(\omega))$ measures the overall utility as the average of individual utilities; this takes into account the total but not the inequality of utilities —i.e. a welfare of $W_U(\omega) = 0.5$ may be achieved with individual utilities 0.5 and 0.5 or with 0.9 and 0.1. The egalitarian welfare of an agreement $W_E(\omega) = min_{A_i \in \mathcal{A}} U_i(\omega)$, takes into account the level of inequality by defining welfare as the minimum utility of the two agents. Since this welfare is very strict, we will use a combination of both called the group welfare $W_G(\omega) = \frac{1}{2}(W_U(\omega)+W_E(\omega))$. In the experiments, group welfare is computed using the ingredient-based utility (hidden to the agents).

The experiments are made for different values of $\alpha \in [0, 1]$ —recall that α is the weight used in the Proposal Selection Heuristic. Thus, when $\alpha = 0$ the agent behaves as an egoist, since all proposed agreements take only into account its individual utility. The higher the α the less egoist is an agent, since it will propose agreements that have less utility for itself but are more similar to the

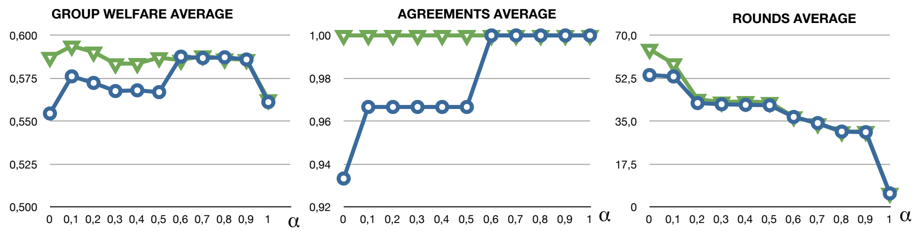

Fig. 1. Group welfare average (left), percentage of times reaching an agreement (center), and number of proposals exchanged (right) for different α values when there is a time limit of 150 rounds (◯) or no limit (∇)

proposals of the other agent. Consequently, the higher the α the faster the aspiration level of the agent decreases during the deliberation process.

Figure 1 shows the open-minded agent relationship with time-sensitiveness. In this experiment, the abscissae are different values of α (for one agent) while the other agent has a random value $\alpha \in [0, 1]$ at each run. The ordinates plot the averages of group welfare, the percent of times an agreement is reached, and the number of rounds needed to reach an agreement. If there is no time limit, the deliberation can spend a lot of time examining a large number of proposals. However, when there is a time limit, being egoist (having values of α close to 0) is a bad option (as shown in Fig. 1): (1) when α increases then the percentage of times in which an agreement is reached also increases (Fig. 1 center), which explains (2) when α is closer to zero group welfare is lower (Fig. 1 left). Finally, the number of exchanged proposals needed to reach an agreement decreases when α increases (Fig. 1 right). Moreover, this last plot shows that the E_i function is useful in estimating the proposals that the other agent may consider satisfactory. For comparison, the centralized method[5] gives a welfare average of $W_G = 0,6735$.

The effect of egoism ($\alpha = 0$) vs. benevolence ($\alpha = 1$) is shown in Fig. 2 for open-minded and narrow-minded agents with a limit of 150 rounds. In this experiment, as well as the following ones, both agents have the same α shown in the abscissa, and both agents are either open-minded or narrow-minded. First, we see in Fig. 2 that open-minded agents achieve agreements with higher welfare values (left), more agreements (center) and with less rounds of deliberation (right) than the close-minded agents. When α is low, both agents are egoistic and thus their aspiration level decrease very slowly, which results in a lower

[5] The centralized aggregation method has complete knowledge. Specifically, this method receives all the cases from the two agents, receives the designer-level ingredient-based evaluation function of each agent and computes the ingredient-based utility for each recipe for each agent. Then performs exhaustive search to find the pair of courses that maximize the group welfare function W_G. This method has complete knowledge of the utility of all alternatives in the experiment and unlimited time to search all combinations, and gives an estimate of the best possible agreement in ideal conditions.

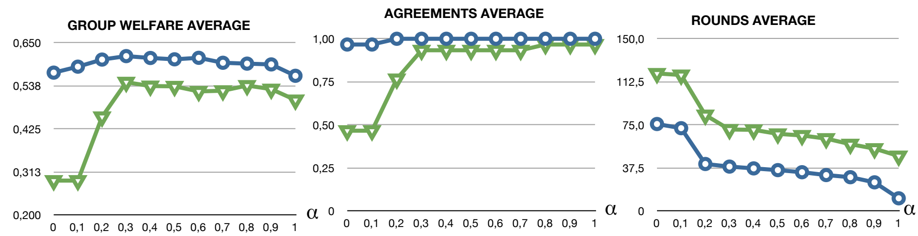

Fig. 2. Group welfare average (left), percentage of times reaching an agreement (center), and number of proposals exchanged (right), depending on the egoist ($\alpha = 0$) vs benevolent ($\alpha = 1$) spectrum for open-minded (◯) and narrow-minded strategies (∇)

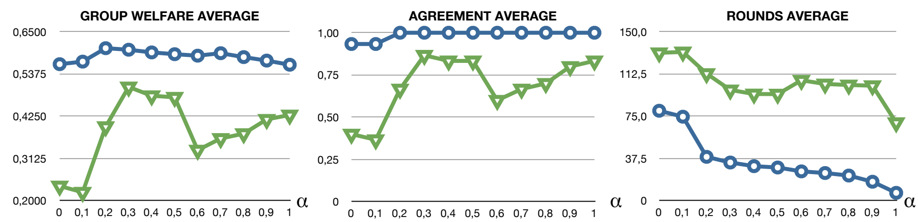

Fig. 3. Group welfare average (left), percentage of times reaching an agreement (center), and number of proposals exchanged (right) for open-minded (◯) and narrow-minded strategies (∇) agents with a 16% of shared recipes

number of agreements reached before the time limit (specially for narrow-minded). When α is higher then the narrow-minded agents improve, but they are still worse than the open-minded ones.

Another parameter affecting the outcome is the size of the case bases for an issue k and the size of the shared recipes. The experiment shown in Fig. 3 has two agents with 150 recipes each agent, a 16% of shared recipes and $M = 150$ rounds, showing the difference between open-minded and narrow-minded agents. In this setting, with smaller case bases and a smaller set of shared recipes, reaching an agreement is more difficult for the narrow-minded agents: group welfare and number of agreements decrease. Moreover, the number of rounds narrow-minded agents need to reach an agreement also increases. The open-minded agents, however, even now that the size of shared alternatives is smaller, keep a similar performance in group welfare, number of agreements, and number of rounds as before.

8 Conclusion

Case-based reasoning is a methodology that allows to address AI under incomplete knowledge by exploiting dynamically available knowledge in a more flexible way. We have addressed here the issues involved in group decisions (modeled here as agreements). Classical mathematical models in social choice assume conditions like perfect knowledge. Weakening this assumption requires a new approach, and we have shown that CBR can deal with incomplete knowledge by exploiting the dynamic exchange of information during deliberation.

Essentially, the classical approach encodes all knowledge in a utility function over known alternatives. We have shown that incorporating a similarity function over the space of possible descriptions of alternatives enables a CBR agent to cope with unknown alternatives (and thus rational ignorance). We have also included a process of deliberation, previous to the group decision itself, that allows to introduce new, unknown alternatives in an incremental way. This approach has been evaluated on the CCC dataset but could be applied to other domains where a suitable similarity over alternatives can be designed.

Future work will address group decision for more than two agents (multilateral agreement), which is a problem whose high complexity is well known. The main idea will be to use the deliberation process to increase the knowledge available to the agents, and then vote. That is to say, the goal is not to achieve an agreement by consensus, since the complexity of the problem would require a very long process. The goal will be that the deliberation process helps the agents gaining new alternatives and acquiring a better model the preferences of other agents before the final decision-making step of voting.

Acknowledgements. This research was partially supported by projects Agreement Technologies (CSD2007-0022) and Next-CBR (TIN2009-13692-C03-01).

References

1. Arrow, K.J.: Social Choice and Individual Values. John Wiley, Chichester (1951)
2. Atkinson, K., Bench-Capon, T., Mcburney, P.: Parmenides: Facilitating deliberation in democracies. Artificial Intelligence and Law 14(4), 261–275 (2006)
3. Chevaleyre, Y., Endriss, U., Lang, J., Maudet, N.: A short introduction to computational social choice. In: van Leeuwen, J., Italiano, G.F., van der Hoek, W., Meinel, C., Sack, H., Plášil, F. (eds.) SOFSEM 2007. LNCS, vol. 4362, pp. 51–69. Springer, Heidelberg (2007)
4. Dryzek, J.S., List, C.: Social choice theory and deliberative democracy: a reconciliation. British Journal of Political Science 33, 1–28 (2003)
5. Elster, J.: Deliberative Democracy (Cambridge Studies in the Theory of Democracy), pp. 11–18. Cambridge University Press, Cambridge (1998)
6. Lang, J.: Logical preference representation and combinatorial vote. Annals of Mathematics and Artificial Intelligence 42, 37–71 (2004)
7. Ontañón, S., Plaza, E.: Case-based learning from proactive communication. In: Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI 2007), pp. 999–1004. IJCAI Press (2007)
8. Ontañón, S., Plaza, E.: Learning and joint deliberation through argumentation in multi-agent systems. In: Proc. AAMAS 2007, pp. 971–978. ACM, New York (2007)
9. Plaza, E., Ontañón, S.: Ensemble case-based reasoning: Collaboration policies for multiagent cooperative CBR. In: Watson, I., Yang, Q. (eds.) ICCBR 2001. LNCS (LNAI), vol. 2080, pp. 437–451. Springer, Heidelberg (2001)